Shortest Path in a Polygon using Sublinear Space*

Sariel Har-Peled†

December 1, 2015

arXiv:1412.0779v2 [cs.CG] 28 Nov 2015

Abstract

We resolve an open problem due to Tetsuo Asano, showing how to compute the shortest path in a polygon, given in a read only memory, using sublinear space and subquadratic time. Specifically, given a simple polygon P with n vertices in a read only memory, and additional working memory of size m, the new algorithm computes the shortest path (in P) in $O(n^2/m)$ expected time, assuming $m = O(n/\log^2 n)$. This requires several new tools, which we believe to be of independent interest.

Specifically, we show that violator space problems, an abstraction of low dimensional linear-programming (and LP-type problems), can be solved using constant space and expected linear time, by modifying Seidel's linear programming algorithm and using pseudo-random sequences.

1. Introduction

Space might not be the final frontier in the design of algorithms but it is an important constraint. Of special interest are algorithms that use sublinear space. Such algorithms arise naturally in streaming settings, or when the data set is massive, and only a few passes on the data are desirable. Another such setting is when one has a relatively weak embedded processor with limited high quality memory. For example, in 2014, flash memory could withstand around 100,000 rewrites before starting to deteriorate. Specifically, imagine a hybrid system having a relatively large flash memory, with significantly smaller RAM. That is to a limited extent the setting in a typical smart-phone[①].

The model. The input is provided in a read only memory, and it is of size n. We have $O(m)$ available space which is a read/write space (i.e., the *work space*). We assume, as usual, that every memory cell is a word, and such a word is large enough to store a number or a pointer. We also assume that the input is given in a reasonable representation[②].

Since a memory cell has $\Omega(\log n)$ bits, for $m = O(1)$, this is roughly the log-space model in complexity. An example of such algorithms are the standard NPComplete reductions, which can all be done in this model. Algorithms developed in this model of limited work memory include: (i) median selection

*Work on this paper was partially supported by NSF AF awards CCF-1421231 and CCF-1217462. A preliminary version of this paper appeared in SoCG 2015 [Har15].

†Department of Computer Science; University of Illinois; 201 N. Goodwin Avenue; Urbana, IL, 61801, USA; sariel@illinois.edu; http://sarielhp.org/.

①For example, a typical smart-phone in 2014 has 2GB of RAM and 16GB of flash memory. I am sure these numbers would be laughable in a few years. So it goes.

②In some rare cases, the "right" input representation can lead directly to sublinear time algorithms. See the work by Chazelle *et al.* [CLM05].

[MR96], (ii) deleting a connected component in a binary image [Asa12], (iii) triangulating a set of points in the plane, computing their Voronoi diagram, and Euclidean minimum spanning tree [Asa+11], to name a few. For more details, see the introduction of Asano *et al.* [Asa+13; Asa+14].

The problem. We are given a simple polygon P with n vertices in the plane, and two points $s, t \in P$ – all provided in a read-only memory. We also have $O(m)$ additional read-write memory (i.e., work space). The task is to compute the shortest path from s to t inside P.

Asano *et al.* [Asa+13] showed how to solve this problem, in $O(n^2/m)$ time, using $O(m)$ space. The catch is that their solution requires quadratic time preprocessing. In a talk by Tetsuo Asano, given in a workshop in honor of his 65th birthday (during SoCG 2014), he posed the open problem of whether this quadratic preprocessing penalty can be avoided. This work provides a positive answer to this question.

If linear space is available. The standard algorithm [LP84] for computing the shortest path in a polygon, triangulates the polygon, (conceptually) computes the dual graph of the triangulation, which yields a tree, with a unique path between the triangles that contains the source s, and the target t. This path specifies the sequence of diagonals crossed by the shortest path, and it is now relatively easy to walk through this sequence of triangles and maintain the shortest paths from the source to the two endpoints of each diagonal. These paths share a prefix path, and then diverge into two concave chains (known together as a *funnel*). Once arriving to the destination, one computes the unique tangent from the destination t to one of these chains, and the (unique) path, formed by the prefix together with the tangent, defines the shortest path, which can be now extracted in linear time. See figure on the right.

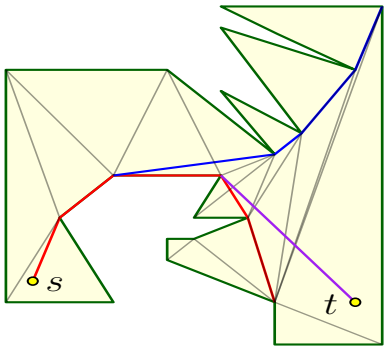

Sketch of the new algorithm. The basic idea is to decompose the polygon into bigger pieces than triangles. Specifically, we break the polygon into canonical pieces each of size $O(m)$. To this end, we break the given polygon P into $\lceil n/m \rceil$ polygonal chains, each with at most m edges. We refer to such a chain as a *curve*. We next use the notion of *corridor decomposition*, introduced[3] by the author [Har14], to (conceptually) decompose the polygon into canonical pieces (i.e., corridors). Oversimplifying somewhat, each edge of the medial axis [Blu67] corresponds to a corridor, which is a polygon having portions of two of the input curves as floor and ceiling, and additional two diagonals of P as gates. It is relatively easy, using constant space and linear time, to figure out for such a diagonal if it separates the source from the destination. Now, start from the corridor containing the source, and figure out which of its two gates the shortest path goes through. We follow this door to the next corridor and continue in this fashion till we reach the destination. Assuming that computing the next corridor can be done in roughly linear time, this algorithm solves the shortest path problem in $O(n^2/m)$ time, as walking through a corridor takes (roughly) linear time, and there are $O(n/m)$ pieces the shortest path might go through. (One also needs to keep track of the funnel being constructed during this walk, and prune parts of it away because of space considerations.)

[3] Many somewhat similar decomposition ideas can be found in the literature. For example, the decomposition of a polygon into monotone polygons so that the pieces, and thus the original polygon, can be triangulated [Ber+08]. Closer to our settings, Kapoor and Maheshwari [KM88] use the dual graph of a triangulation to create a decomposition of the free space among polygonal obstacles into corridors. Nevertheless, the decomposition scheme of [Har14] is right for our nefarious purposes, but the author would not be surprised if it was known before. Well, at least this footnote is new!

point location queries in a canonical decomposition. To implement the above, we need a way to perform a point location query in the corridor decomposition, without computing it explicitly. Specifically, we have an implicit decomposition of the polygon into subpolygons, and we walk through this decomposition by performing a sequence of point location queries on the shared borders between these pieces.

More generally, we are interested in any canonical decomposition that partition the underlying space into cells. Such a partition is induced by a set of objects, and every cell is defined by a constant number of objects. Standard examples of such partitions are (i) vertical decomposition of segments in the plane, or (ii) bottom vertex triangulation of the Voronoi diagram of points in $\mathbb{R}^3$. Roughly speaking, any partition that complies with the Clarkson-Shor framework [Cla88; CS89] is such a canonical decomposition, see Section 2.1.2 for details.

If space and time were not a constraint, we could build the decomposition explicitly, Then a standard point location query in the history DAG would yield the desired cell. Alternatively, one can perform this point location query in the history DAG implicitly, without building the DAG beforehand, but it is not obvious how to do so with limited space. Surprisingly, at least for the author, this task can be solved using techniques related to low-dimensional linear programming.

Violator spaces. Low dimensional linear programming can be solved in linear time [Meg84]. Sharir and Welzl [SW92] introduced *LP-type* problems, which are an extension of linear programming. Intuitively, but somewhat incorrectly, one can think about LP-Type algorithms as solving low-dimensional convex programming, although Sharir and Welzl [SW92] used it to decide in linear time if a set of axis-parallel rectangles can be stabbed by three points (this is quite surprising as this problem has no convex programming flavor). LP-type problems have the same notions as linear programming of bases, and an objective function. The function scores such bases, and the purpose is to find the basis that does not violate any constraint and minimizes (or maximizes) this objective. A natural question is how to solve such problems if there is no scoring function of the bases.

This is captured by the notion of *violator spaces* [Rüs07; Ško07; Gär+06; Gär+08; BG11]. The basic idea is that every subset of constraints is mapped to a unique basis, every basis has size at most δ (δ is the dimension of the problem, and is conceptually a constant), and certain conditions on consistency and monotonicity hold. Computing the basis of a violator space is not as easy as solving LP-type problems, because without a clear notion of progress, one can cycle through bases (which is not possible for LP-type problems). See Šavroň [Ško07] for an example of such cycling. Nevertheless, Clarkson's algorithm [Cla95] works for violator spaces [BG11].

We revisit the violator space framework, and show the following:

(A) Because of the cycling mentioned above, the standard version of Seidel's linear programming algorithm [Sei91] does not work for violator spaces. However, it turns out that a variant of Seidel's algorithm does work for violator spaces.

(B) We demonstrate that violator spaces can be used to solve the problem of point location in canonical decomposition. While in some cases this point location problem can be stated as an LP-type problem, stating it as a violator space problem seems to be more natural and elegant.

(C) The advantages of Seidel's algorithm is that except for constant work space, the only additional space it needs is to store the random permutations it uses. We show that one can use pseudo-random generators (PRGs) to generate the random permutation, so that there is no need to store it explicitly. This is of course well known – but the previous analysis [Mul94] for linear programming implied only that the expected running time is $O(n \log^{\delta-1} n)$, where δ is the combinatorial dimension. Building on Mulmuley's work [Mul94], we do a somewhat more careful analysis, showing that in this case

one can use backward analysis on the random ordering of constraints generated, and as such the expected running time remains linear.

This implies that one can solve violator space problems (and thus, LP and LP-type problems) in constant dimension, using constant space, in expected linear time.

This compares favorably with previous work: (A) The algorithm for LP by Mulmuley [Mul94], mentioned above, that uses constant space, and takes $O(n \log^{\delta-1} n)$ expected time. (B) Chan and Chen's [CC07, Theorem 3.12] algorithm for LP and LP-type problems that uses $O(\log n)$ space and takes expected linear time.

Paper organization. We present the new algorithm for computing the basis of violator spaces in Section 2. The adaptation of the algorithm to work with constant space is described in Section 2.3. We describe corridor decomposition and its adaptation to our setting in Section 3. We present the shortest path algorithm in Section 4.

2. Violator spaces and constant space algorithms

First, we review the formal definition of *violator spaces* [Rüs07; Ško07; Gär+06; Gär+08; BG11]. We then show that a variant of Seidel's algorithm for linear programming works for this abstract setting, and show how to adapt it to work with constant space and in expected linear time.

2.1. Formal definition of violator space

Before delving into the abstract framework, let us consider the following concrete example – hopefully it would help the reader in keeping track of the abstraction.

Example 2.1. We have a set H of n segments in the plane, and we would like to compute the vertical trapezoid of $\mathcal{A}^{|}(H)$ that contains, say, the origin, where $\mathcal{A}^{|}(H)$ denote the vertical decomposition of the arrangement formed by the segments of H. For a subset $X \subseteq H$, let $\tau(X)$ be the vertical trapezoid in $\mathcal{A}^{|}(X)$ that contains the origin. The vertical trapezoid $\tau(X)$ is defined by at most four segments, which are the *basis* of X. A segment $f \in H$ *violates* $\tau = \tau(X)$, if it intersects the interior of $\tau(X)$. The set of segments of H that intersect the interior of τ, denoted by $\mathrm{cl}(\tau)$ or $\mathrm{cl}(X)$, is the *conflict list* of τ.

Somewhat informally, a violator space identifies a vertical trapezoid $\tau = \tau(X)$, by its conflict list $\mathrm{cl}(X)$, and not by its geometric realization (i.e., τ).

Definition 2.2. A ***violator space*** is a pair $\mathcal{V} = (H, \mathrm{cl})$, where H is a finite set of ***constraints***, and $\mathrm{cl} : 2^H \to 2^H$ is a function, such that:

- **Consistency**: For all $X \subseteq H$, we have that $\mathrm{cl}(X) \cap X = \emptyset$.
- **Locality**: For all $X \subseteq Y \subseteq H$, if $\mathrm{cl}(X) \cap Y = \emptyset$ then $\mathrm{cl}(X) = \mathrm{cl}(Y)$.
- **Monotonicity**: For all $X \subseteq Y \subseteq Z \subseteq H$, if $\mathrm{cl}(X) = \mathrm{cl}(Z)$ then $\mathrm{cl}(X) = \mathrm{cl}(Y) = \mathrm{cl}(Z)$.

A set $B \subseteq X \subseteq H$ is a ***basis*** of X, if $\mathrm{cl}(B) = \mathrm{cl}(X)$, and for any proper subset $B' \subset B$, we have that $\mathrm{cl}(B') \neq \mathrm{cl}(B)$. The ***combinatorial dimension***, denoted by δ, is the maximum size of a basis.

Note that consistency and locality implies monotonicity. For the sake of concreteness, it is also convenient to assume the following (this is strictly speaking not necessary for the algorithm).

Definition 2.3. For any $X \subseteq H$ there is a unique ***cell*** $\tau(X)$ associated with it, where for any $X, Y \subseteq H$, we have that if $\mathrm{cl}(X) \neq \mathrm{cl}(Y)$ then $\tau(X) \neq \tau(Y)$. Consider any $X \subseteq H$, and any $f \in H$. For $\tau = \tau(X)$, the constraint f ***violates*** τ if $f \in \mathrm{cl}(X)$ (or alternatively, f *violates* X).

Finally, we assume that the following two basic operations are available:

- **violate**(f, B): Given a basis B (or its cell $\tau = \tau(B)$) and a constraint f, it returns true $\iff$ f violates τ.
- **compBasis**(X): Given a set X with at most $(\delta + 1)^2$ constraints, this procedure computes basis(X), where δ is the combinatorial dimension of the violator space[④]. For δ a constant, we assume that this takes constant time.

2.1.1. Linear programming via violator spaces

Consider an instance I of linear programming in $\mathbb{R}^d$ – here the LP is defined by a collection of linear inequalities with d variables. The instance I induces a polytope P in $\mathbb{R}^d$, which is the *feasible domain* – specifically, every inequality induces a halfspace, and their intersection is the polytope.

The following interpretation of the feasible polytope is somewhat convoluted, but serves as a preparation for the next example. The vertices V of the polytope P induce a triangulation (assuming general position) of the sphere of directions, where a direction v belongs to a vertex p, if and only if p is an extreme vertex of P in the direction of v. Now, the objective function of I specifies a direction v_I, and in solving the LP, we are looking for the extreme vertex of P in this direction.

Put differently, every subset H of the constraints of I, defines a triangulation $\mathcal{T}(H)$ of the sphere of directions. So, let the cell of H, denoted by $\tau = \tau(H)$, be the spherical triangle in this decomposition that contains v_I. The basis of H is the subset of constraints that define $\tau(H)$. A constraint f of the LP violates τ if the vertex induced by the basis basis(H) (in the original space), is on the wrong side of f.

Thus solving the LP instance $I = (H, v_I)$ is no more than performing a point location query in the spherical triangulation $\mathcal{T}(H)$, for the spherical triangle that contains v_I.

Remark. An alternative, and somewhat more standard, way to see this connection is via geometric duality [Har11, Chapter 25] (which is *not* LP duality). The duality maps the upper envelope of hyperplanes to a convex-hull in the dual space (i.e., we assume here that in the given LP all the halfspaces (i.e., each constraint corresponds to a halfspace) contain the positive ray of the x_d-axis – otherwise, we need to break the constraints of the LP into two separate families, and apply this reduction separately to each family). Then, extremal query on the feasible region of the LP, in the dual, becomes a vertical ray shooting query on the upper portion of the convex-hull of the dual points – that is, a point location query in the projection of the triangulation of the upper portion of the convex-hull of the dual points.

2.1.2. Point location via violator spaces

Example 2.1 hints to a more general setup. So consider a space decomposition into canonical cells induced by a set of objects. For example, segments in the plane, with the canonical cells being the vertical trapezoids. More generally, consider any decomposition of a domain into simple canonical cells induced by objects, which falls under the Clarkson-Shor framework [Cla88; CS89] (defined formally below). Examples of this include point location in a (i) Delaunay triangulation, (ii) bottom vertex

[④] We consider basis(X) to be unique (that is, we assume implicitly that the input is in general position). This can be enforced by using lexicographical ordering, if necessary, among the defining bases always using the lexicographically minimum basis.

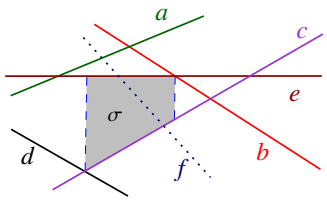

Figure 2.1: Vertical decomposition and the Clarkson-Shor framework. The defining set of σ is $D(\sigma) = \{a, b, c, d\}$, and its stopping set is $K(\sigma) = \{f\}$.

triangulation in an arrangement of hyperplanes, (iii) and vertical decomposition of arrangements of curves and surfaces in two or three dimensions, to name a few.

Definition 2.5 (Clarkson-Shor framework [Cla88; CS89]). Let $\mathcal{D}$ be an underlying domain, and let H be a set of objects, such that any subset $R \subseteq H$ decomposes $\mathcal{D}$ into canonical cells $\mathcal{F}(R)$. This decomposition complies with the Clarkson-Shor framework, if we have that for every cell τ that arises from such a decomposition has a *defining set* $D(\tau) \subseteq H$. The size of such a defining set is assumed to be bounded by a constant δ. The *stopping set* (or conflict list) $K(\tau)$ of τ is the set of objects of H such that including any object of $K(\tau)$ in R prevents τ from appearing in $\mathcal{F}(R)$. We require that for any $R \subseteq H$, the following conditions hold:

(i) For any $\tau \in \mathcal{F}(R)$, we have $D(\tau) \subseteq R$ and $R \cap K(\tau) = \emptyset$.
(ii) If $D(\tau) \subseteq R$ and $K(\tau) \cap R = \emptyset$, then $\tau \in \mathcal{F}(R)$.

For a detailed discussion of the Clarkson-Shor framework, see Har-Peled [Har11, Chapter 8]. We next provide a quick example for the reader unfamiliar with this framework.

Example 2.6. Consider a set of segments H in the plane. For a subset $R \subseteq H$, let $\mathcal{A}^{|}(R)$ denote the ***vertical decomposition*** of the plane formed by the arrangement $\mathcal{A}(R)$ of the segments of R. This is the partition of the plane into interior disjoint vertical trapezoids formed by erecting vertical walls through each vertex of $\mathcal{A}(R)$. Here, each object is a segment, a region is a vertical trapezoid, and $\mathcal{F}(R)$ is the set of vertical trapezoids in $\mathcal{A}^{|}(R)$. Each trapezoid $\sigma \in \mathcal{F}(R)$ is defined by at most four segments (or lines) of R that define the region covered by the trapezoid σ, and this set of segments is $D(\sigma)$. Here, $K(\sigma)$ is the set of segments of H intersecting the interior of the trapezoid σ (see Figure 2.1).

Lemma 2.7. *Consider a canonical decomposition of a domain into simple cells, induced by a set of objects, that complies with the Clarkson-Shor framework [Cla88; CS89]. Then, performing a point location query in such a domain is equivalent to computing a basis of a violator space.*

Proof: This follows readily from definition, but we include the details for the sake of completeness. We use the notations of Definition 2.5 above.

So consider a fixed point $p \in \mathcal{D}$, and the task at hand is to compute the cell $\tau \in \mathcal{F}(H)$ that contains p (here, assuming general position implies that there is a unique such cell). In particular, for $R \subseteq H$, we define $B = \text{basis}(R)$ to be the defining set of the cell τ of $\mathcal{F}(R)$ that contains p, and the conflict list to be $\text{cl}(R) = K(\tau)$.

We claim that computing $\text{basis}(H)$ is a violator space problem. We need to verify the conditions of Definition 2.2, which is satisfyingly easy. Indeed, consistency is condition (i), and locality is condition (ii) in Definition 2.5. (We remind the reader that consistency and locality implies monotonicity, and thus we do not have to verify that it holds.)

```
solveVS(W, X):
    ⟨f_1, ..., f_m⟩: A random permutation of the constraints of X.
    B_0 ← compBasis(W)
    for i = 1 to m do
        if violate(f_i, B_{i-1}) then
            B_i ← solveVS(W ∪ B_{i-1} ∪ {f_i}, {f_1, ..., f_i})
        else
            B_i ← B_{i-1}
    return B_m
```

Figure 2.2: The algorithm for solving violator space problems. The parameter W is a set of $O(\delta^2)$ witness constraints, and X is a set of m constraints. The function return basis$(W \cup X)$. To solve a given violator space, defined implicitly by the set of constraints H, and the functions **violate** and **compBasis**, one calls **solveVS**$(\{\}, H)$.

It seems that for all of these point location problems, one can solve them directly using the LP-type technique. However, stating these problems as violator space instances is more natural as it avoids the need to explicitly define an artificial ordering over the bases, which can be quite tedious and not immediate (see Appendix B for an example).

2.2. The algorithm for computing the basis of a violator space

The input is a violator space $\mathcal{V} = (H, \mathrm{cl})$ with $n = |H|$ constraints, having combinatorial dimension δ.

2.2.1. Description of the algorithm

The algorithm is a variant of Seidel's algorithm [Sei91] – it picks a random permutation of the constraints, and computes recursively in a randomized incremental fashion the basis of the solution for the first i constraints. Specifically, if the ith constraint violates the basis B_{i-1} computed for the first $i-1$ constraints, it calls recursively, adding the constraints of B_{i-1} and the ith constraint to the set of constraints that must be included whenever computing a basis (in the recursive calls). The resulting code is depicted in Figure 2.2.

The only difference with the original algorithm of Seidel, is that the recursive call gets the set $W \cup B_{i-1} \cup \{f_i\}$ instead of basis$(B_{i-1} \cup \{f_i\})$ (which is a smaller set). This modification is required because of the potential cycling between bases in a violator space.

2.2.2. Analysis

The key observation is that the depth of the recursion of **solveVS** is bounded by δ, where δ is the combinatorial dimension of the violator space. Indeed, if f_i violates a basis, the constraints added to the witness set W guarantee that any subsequent basis computed in the recursive call contains f_i, as testified by the following lemma.

Lemma 2.8. *Consider any set $X \subseteq H$. Let $B = \mathrm{basis}(X)$, and let f be a constraint in $H \setminus X$ that violates B. Then, for any subset Y such that $B \cup \{f\} \subseteq Y \subseteq X \cup \{f\}$, we have that $f \in \mathrm{basis}(Y)$.*

Proof: Assume that this is false, and let Y be the bad set with $B' = \text{basis}(Y)$, such that $f \notin B'$. Since $f \in Y$, by consistency, $f \notin \text{cl}(Y)$, see Definition 2.2. By definition $\text{cl}(Y) = \text{cl}(B')$, which implies that $f \notin \text{cl}(B')$; that is, f does not violate B'.

Now, by monotonicity, we have $\text{cl}(Y) = \text{cl}(Y \setminus \{f\}) = \text{cl}(B')$.

By assumption, $B \subseteq Y \setminus \{f\}$, which implies, again by monotonicity, as $B \subseteq Y \setminus \{f\} \subseteq X$, that $\text{cl}(X) = \text{cl}(Y \setminus \{f\}) = \text{cl}(B)$, as $B = \text{basis}(X)$. But that implies that $\text{cl}(B) = \text{cl}(Y \setminus \{f\}) = \text{cl}(B')$. As $f \notin \text{cl}(Y)$, this implies that f does not violate B, which is a contradiction. ∎

Lemma 2.9. *The depth of the recursion of* **solveVS**, *see Figure 2.2*$_{p7}$, *is at most δ, where δ is the combinatorial dimension of the given instance.*

Proof: Consider a sequence of k recursive calls, with $W_0 \subseteq W_1 \subseteq W_2 \subseteq \cdots \subseteq W_k$ as the different values of the parameter W of **solveVS**, where $W_0 = \emptyset$ is the value in the top-level call. Let g_j, for $j = 1, \ldots, k$, be the constraint whose violation triggered the jth level call. Observe that $g_j \in W_j$, and as such all these constraints must be distinct (by consistency). Furthermore, we also included the basis B'_j, that g_j violates, in the witness set W_j. As such, we have that $W_j = \{g_1, \ldots, g_j\} \cup B'_1 \cup \cdots \cup B'_j$. By Lemma 2.8, in any basis computation done inside the recursive call **solveVS**$(W_j, \ldots)$, it must be that $g_j \in \text{basis}(W_t)$, for any $t \geq j$. As such, we have $g_1, \ldots, g_k \in \text{basis}(W_k)$. Since a basis can have at most δ elements, this is possible only if $k \leq \delta$, as claimed. ∎

Theorem 2.10. *Given an instance of violator space $\mathcal{V} = (H, \text{cl})$ with n constraints, and combinatorial dimension δ, the algorithm* **solveVS**$(\emptyset, H)$, *see Figure 2.2, computes* $\text{basis}(H)$. *The expected number of violation tests performed is bounded by $O(\delta^{\delta+1} n)$. Furthermore, the algorithm performs in expectation $O\Big((\delta \ln n)^{\delta}\Big)$ basis computations (on sets of constraints that contain at most $\delta(\delta+1)$ constraints).*

In particular, for constant combinatorial dimension δ, with violation test and basis computation that takes constant time, this algorithm runs in $O(n)$ expected time.

Proof: Lemma 2.9 implies that the recursion tree has bounded depth, and as such this algorithm must terminate. The correctness of the result follows by induction on the depth of the recursion. By Lemma 2.9, any call of depth δ, cannot find any violated constraint in its subproblem, which means that the returned basis is indeed the basis of the constraints specified in its subproblem. Now, consider a recursive call at depth $j < \delta$, which returns a basis B_m when called on the constraints $\langle f_1, \ldots, f_m \rangle$. The basis B_m was computed by a recursive call on some prefix $\langle f_1, \ldots, f_i \rangle$, which was correct by (reverse) induction on the depth, and none of the constraints $f_{i+1}, \ldots, f_m$ violates B_m, which implies that the returned basis is a basis for the given subproblem. Thus, the result returned by the algorithm is correct.

Let $C_k(m)$ be the expected number of basis computations performed by the algorithm when run at recursion depth k, with m constraints. We have that $C_\delta(m) = 1$, and $C_k(m) = 1 + \mathbf{E}\Big[\sum_{i=1}^{m} X_i C_{k+1}(i)\Big]$, where X_i is an indicator variable that is one, if and only if the insertion of the ith constraint caused a recursive call. We have, by backward analysis, that $\mathbf{E}[X_i] = \mathbf{Pr}[X_i = 1] \leq \min(\delta/i, 1)$. As such, by linearity of expectations, we have $C_k(m) \leq 1 + \sum_{i=1}^{m} \min\big(\frac{\delta}{i}, 1\big) C_{k+1}(i)$. As such, we have

$$C_{\delta-1}(m) \leq 1 + \sum_{i=1}^{m} \min\left(\frac{\delta}{i}, 1\right) C_\delta(i) = \delta + \sum_{i=\delta}^{m} \frac{\delta}{i} \leq \delta + \delta \ln(m/\delta) \leq \delta \ln(m),$$

assuming $\delta \geq e$. We conclude that $C_{\delta-1}(m) \leq \delta \ln(m)$ and $C_0(m) = O\Big((\delta \ln m)^{\delta}\Big)$.

As for the expected number of violation tests, a similar analysis shows that $V_\delta(m) = m$ and $V_{\delta-j}(m) = m + \sum_{i=1}^{m} \min\left(\frac{\delta}{i}, 1\right) V_{\delta-(j-1)}(i)$. As such, assuming inductively that $V_{\delta-(j-1)}(m) \leq j\delta^{j-1}m$, we have that

$$V_{\delta-j}(m) \leq m + \sum_{i=1}^{m} \frac{\delta}{i} j\delta^{j-1} i \leq m + j\delta^j m \leq (j+1)\delta^j m,$$

implying that $V_0(m) = O(\delta^{\delta+1}m)$, which also bounds the running time. ∎

Remark 2.11. While the constants in Theorem 2.10 are not pretty, we emphasize that the O notation in the bounds do not hide constants that depends on δ.

2.3. Solving the violator space problem with constant space and linear time

The idea for turning **solveVS** into an algorithm that uses little space, is observing that the only thing we need to store (implicitly) is the random permutation used by **solveVS**.

2.3.1. Generating a random permutation using pseudo-random generators

To avoid storing the permutation, one can use pseudo-random techniques to compute the permutation on the fly. For our algorithm, we do not need a permutation - any random sequence that has uniform distribution over the constraints and is sufficiently long, would work.

Lemma 2.12. *Fix an integer $\phi > 0$, a prime integer n, and an integer constant $c' \geq 12$. One can compute a random sequence of numbers $X_1, \ldots, X_{c'n} \in [\![n]\!] = \{1, \ldots, n\}$, such that:*

(A) The probability of $X_i = j$ is $1/n$, for any $i \in [\![n]\!]$ and $j \in [\![c'n]\!]$.

(B) The sequence is ϕ-wise independent.

(C) Using $O(c'\phi)$ space, given an index i, one can compute X_i in $O(\phi)$ time.

Proof: This is a standard pseudo-random generator (PRG) technique, described in detail by Mulmuley [Mul94, p. 399]. We outline the idea. Randomly pick $\phi + 1$ coefficients $\alpha_0, \ldots, \alpha_\phi \in \{0, \ldots, n-1\}$ (uniformly and independently), and consider the random polynomial $f(x) = \sum_{i=0}^{\phi} \alpha_i x^i$, and set $p(x) = (f(x) \mod n)$. Now, set $X_i = 1 + p(i)$, for $i = 1, \ldots, n$. It is easy to verify that the desired properties hold. To extent this sequence to be of the desired length, pick randomly c' such polynomials, and append their sequences together to get the desired longer sequence. It is easy to verify that the longer sequence is still ϕ-wise independent. ∎

We need the following concentration result on $2k$-wise independent sequences – for the sake completeness, we provide a proof at Appendix A.1.

Lemma 2.13 (Theorem A.2.1, [Mul94]). *Let $X_1, \ldots, X_n$ be n random indicator variables that are 2ψ-wise independent, where $p = \mathbf{Pr}[X_i = 1] \geq 1/n$, for all i. Let $Y = \sum_i X_i$, and let $\mu = \mathbf{E}[Y] = pn$. Then, we have that $\mathbf{Pr}\left[Y = 0\right] \leq \mathbf{Pr}\left[|Y - \mu| \geq \mu\right] \leq \left(4\psi^2/\mu\right)^\psi / \sqrt{2\psi}$.*

The following lemma testifies that this PRG sequence, with good probability, contains the desired basis (as such, conceptually, we can think about it as being a permutation of $[\![n]\!]$).

Lemma 2.14. *Let $B \subseteq [\![n]\!]$ be a specific set of δ numbers. Fix an integer $\phi \geq 8 + 2\delta$, and let $\mathcal{X} = \langle X_1, \ldots, X_{c'n} \rangle$ be a ϕ-wise independent random sequence of numbers, each uniformly distributed in $[\![n]\!]$, where c' is any constant $\geq 8(5 + \lceil \ln \delta \rceil)^2$. Then, the probability that the elements of B do not appear in $\mathcal{X}$ is bounded by, say, $1/20$.*

Proof: Fix an element b of B. Let Y_i be an indicator variable that is one $\iff$ $X_i = b$. Observe that $p = \mathbf{Pr}[Y_i = 1] = 1/n$. As such, we have that $\mu = \mathbf{E}\left[\sum_{i=1}^{c'n} Y_i\right] = c' = 8k^2$, for $k = 5 + \lceil \ln \delta \rceil$.

As $k \leq \phi/2$, we can interpret $\mathcal{X}$ as a $2k$-wise independent sequence. By Lemma 2.13, we have that the probability that b does not appear in $\mathcal{X}$ is bounded by $\mathbf{Pr}[\sum_i Y_i = 0] \leq (4k^2/\mu)^k/\sqrt{2k}$. We want this probability to be smaller than $1/(20\delta)$, Since $\mu = 8k^2$, this is equivalent to $(4k^2/\mu)^k \delta \leq \sqrt{2k}/20 \iff \delta \leq 2^k\sqrt{2k}/20$ which holds as $2^k \geq 32\delta$.

Now, by the union bound, the probability that any element of B does not appear in $\mathcal{X}$ is bounded by $\delta/(20\delta) = 1/20$. ∎

Remark 2.15. There are several low level technicalities that one needs to address in using such a PRG sequence instead of a truly random permutation:

(A) REPEATED NUMBERS ARE NOT A PROBLEM: the algorithm **solveVS** (see Figure 2.2$_{\text{p7}}$) ignores a constraint that is being inserted for the second time, since it cannot violate the current basis.

(B) VERIFYING THE SOLUTION: The sequence (of the indices) of the constraints used by the algorithm would be $X_1, \ldots, X_{c'n}$. This sequence might miss some constraints that violate the computed solution.
As such, in a second stage, the algorithm checks if any of the constraints $1, 2, \ldots, n$ violates the basis computed. If a violation was found, then the sequence generated failed, and the algorithm restarts from scratch – resetting the PRG used in this level, regenerating the random keys used to initialize it, and rerun it to generate a new sequence. Lemma 2.14 testifies that the probability for such a failure is at most $1/20$ – thus restarting is going to be rare.

(C) INDEPENDENCE BETWEEN LEVELS: We will use a different PRG for each level of the recursion of **solveVS**. Specifically, we generate the keys used in the PRG in the beginning of each recursive call. Since the depth of the recursion is δ, that increases the space requirement by a factor of δ.

(D) IF THE SUBPROBLEM SIZE IS NOT A PRIME: In a recursive call, the number of constraints given (i.e., m) might not be a prime. To this end, the algorithm can store (non-uniformly[5]), a list of primes, such that for any m, there is a prime $m' \geq m$ that is at most twice bigger than m[6]. Then the algorithm generates the sequence modulo m', and ignores numbers that are larger than m. This implies that the sequence might contain invalid numbers, but such numbers are only a constant fraction of the sequence, so ignoring them does not change the running time analysis of our algorithm. (More precisely, this might cause the running time of the algorithm to deteriorate by a factor of $\exp(O(\delta))$, but as we consider δ to be a constant, this does not effect our analysis.)

One needs now to prove that backward analysis still works for our algorithm for violator spaces. The proof of the following lemma is implied by a careful tweaking of Mulmuley's analysis – we provide the details in Appendix A.

[5] That is, sufficiently large sequence of such numbers have to be hard coded into the program, so that it can handle any size input the program is required to handle. This notion is more commonly used in the context of nonuniform circuit complexity.

[6] Namely, the program hard codes a list of such primes. The author wrote a program to compute such a list of primes, and used it to compute 50 primes that cover the range all the way to 10^{15} (the program runs in a few seconds). This would require hard coding $O(\log^2 n)$ bits storing these primes. Alternatively, one can generate such primes on the fly (and no hard coding of bits is required), but the details become more complicated, and this issue is somewhat orthogonal to the main trust of the paper.

Nevertheless, here is a short outline of the idea: Legendre's conjectured that for any integer $x > 0$, there is a prime in the interval $[x, x + O(\sqrt{x})]$ (Cramér conjectured that this interval is of length $O(\log^2 x)$). As such, since one can test primality for a number x in time polynomial in the number of bits encoding x, it follows that one can find a prime close to 2^i in time which is $o(2^i)$, which is all we need for our scheme.

Lemma 2.16 (See Appendix A). *Consider a violator space $\mathcal{V} = (H, \mathrm{cl})$ with $n = |H|$, and combinatorial dimension δ. Let $\mathcal{X} = X_1, X_2, \ldots,$ be a random sequence of constraints of H generated by a ϕ-wise independent distribution (with each X_i having a uniform distribution), where $\phi > 6\delta + 9$ is a constant. Then, for $i > 2\delta$, the probability that X_i violates $B = \mathrm{basis}(X_1, \ldots, X_{i-1})$ is $O(1/i)$.*

Remark. Lemma 2.16 states that backwards analysis works on pseudo-random sequences when applied to the Clarkson-Shor [Cla88; CS89] settings, where the defining set of the event (or basis) of interest has constant size. However, there are cases when one would like to apply backwards analysis when the defining set might have an unbounded size, and the above does not hold in this case.

For an example where a defining set might have unbounded size, consider a random permutation of points, and the event being a point being a vertex of the convex-hull of the points inserted so far. This is used in a recent proof showing that the complexity of the convex-hull of n points, picked uniformly and randomly, in the hypercube in d dimensions has $O(\log^{d-1} n)$ vertices, with high probability [CHR15],

2.3.2. The result

Theorem 2.18. *Given an instance of a violator space $\mathcal{V} = (H, \mathrm{cl})$ with n constraints, and combinatorial dimension δ, one can compute $\mathrm{basis}(H)$ using $O\big(\delta^2 \log^2 \delta\big)$ space. For some constant $\zeta = O\big(\delta \log^2 \delta\big)$, we have that:*
(A) The expected number of basis computations is $O\big((\zeta \ln n)^{\delta}\big)$, each done over $O(\delta^2)$ constraints.
(B) The expected number of violation tests performed is $O\big(\zeta^{\delta} n\big)$.
(C) The expected running time (ignoring the time to do the above operations) is $O\big(\zeta^{\delta} n\big)$.

Proof: The algorithm is described above. As for the analysis, it follows by plugging the bound of Lemma 2.16 into the proof of Theorem 2.10.

The only non-trivial technicality is to handle the case that the PRG sequence fails to contain the basis. Formally, abusing notations somewhat, consider a recursive call on the constraints indexed by $[\![n]\!]$, and let B be the desired basis of the given subproblem. By Lemma 2.14, the probability that B is not contained in the generated PRG is bounded by $1/20$ – and in such a case the sequence has to be regenerated till success. As such, in expectation, this has a penalty factor of (say) 2 on the running time in each level. Overall, the analysis holds with the constants deteriorating by a factor of (at most) 2^{δ}. ■

Remark 2.19. Note, that the above pseudo-random generator technique is well known, but using it for linear programming by itself does not make too much sense. Indeed, pseudo-random generators are sometimes used as a way to reduce the randomness consumed by an algorithm. That in turn is used to derandomize the algorithm. However, for linear programming Megiddo's original algorithm was already linear time deterministic. Furthermore, Chazelle and Matoušek [CM96], using different techniques showed that one can even derandomize Clarkson's algorithm and get a linear running time with a better constant.

Similarly, using PRGs to reduce space of algorithms is by now a standard technique in streaming, see for example the work by Indyk [Ind06], and references therein.

3. Corridor decomposition

3.1. Construction

The decomposition here is similar to the decomposition described by the author in a recent work [Har14].

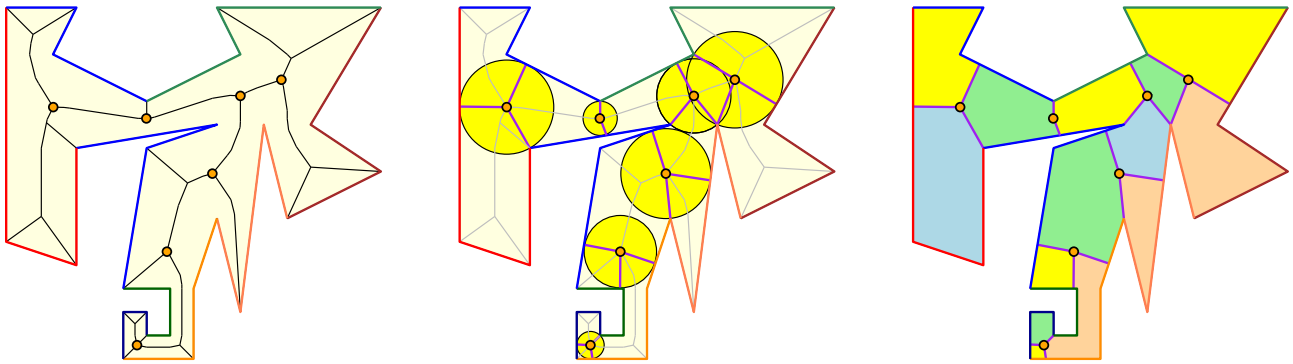

(i) (ii) (iii)

Figure 3.1: An example of a corridor decomposition for a polygon: (i) Input curves, medial axis and active vertices, (ii) their critical circles, and their spokes, and (iii) the resulting corridor decomposition.

Definition 3.1 (Breaking a polygon into curves). Let the given polygon P have the vertices $v_1, \ldots v_n$ in counterclockwise order along its boundary. Let σ_i be the polygonal curve having the vertices $v_{(i-1)m+1}, v_{(i-1)m+2}, v_{(i-1)m+3}, \ldots, v_{im+1}$, for $i = 1, \ldots, \overline{n} - 1$, where

$$\overline{n} = \lfloor (n-1)/m \rfloor + 1.$$

The last polygonal curve is $\sigma_{\overline{n}} = v_{(\overline{n}-1)m+1}, v_{(\overline{n}-1)m+2}, \ldots v_n, v_1$. Note, that given P in a read only memory, one can encode any curve σ_i using $O(1)$ space. Let $\Gamma = \{\sigma_1, \ldots, \sigma_{\overline{n}}\}$ be the resulting set of polygonal curves. From this point on, a ***curve*** refers to a polygonal curve generated by this process.

3.1.1. Corridor decomposition for the whole polygon

Next, consider the medial axis of P restricted to the interior of P (see Chin *et al.* [CSW99] for the definition and an algorithm for computing the medial axis of a polygon). A vertex v of the medial axis corresponds to a disk D, that touches the boundary of P in two or three points (by the general position assumption, that the medial axis of the polygon does not have degenerate vertices, not in any larger number of points). The medial axis has the topological structure of a tree.

To make things somewhat cleaner, we pretend that there is a little hole centered at every vertex of the polygon if it is the common endpoint of two curves. This results in a medial axis edge that comes out of the vertex as an angle bisector, both for an acute angle (where a medial axis edge already exists), and for obtuse angles, see Figure 3.1 and Figure 3.3.

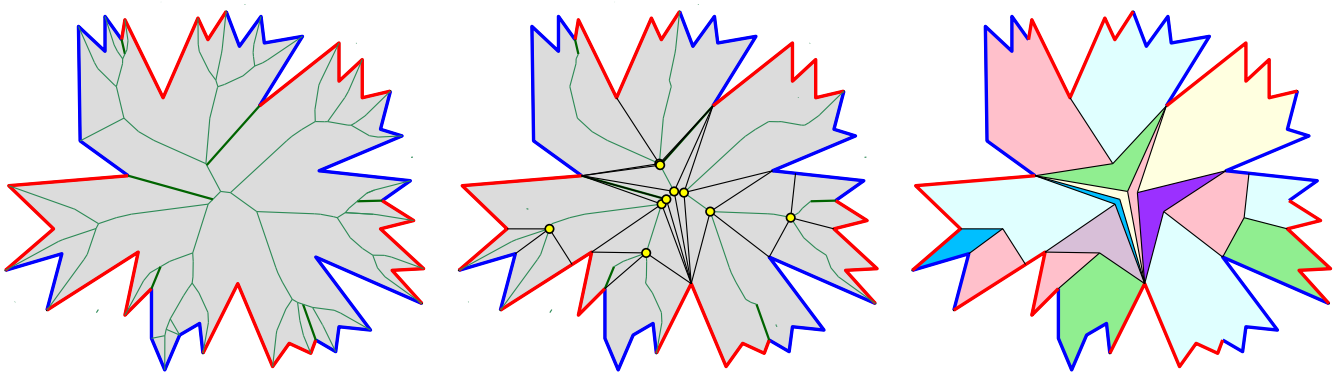

Figure 3.3: Another example of a corridor decomposition for a polygon: (i) Input polygon and its curves and its medial axis (the thick lines are the angle bisectors for the obtuse angles where two curves meet), (ii) active vertices and their spokes (with a reduced medial axis), and (iii) the resulting corridor decomposition.

A vertex of the medial axis is ***active*** if its disk touches three different curves of Γ. It is easy to verify that there are $O(\overline{n})$ active vertices. The segments connecting an active vertex to the three (or two) points of tangency of its empty disk with the boundary of P are its ***spokes*** or ***doors***. Introducing these spokes breaks the polygon into the desired ***corridors***.

Such a corridor is depicted on the right. It has a ***floor*** and a ***ceiling*** – each is a subcurve of some input curve. In addition, each corridor might have four additional doors grouped into "double" doors. Such a double door is defined by an active vertex and two segments emanating from it to the floor and ceiling curves, respectively. We refer to such a double door as an ***gate***. As such, specifying a single corridor requires encoding the vertices of the two gates and the floor and ceiling curves, which can be done in $O(1)$ space, given the two original curves. In particular, a corridor is defined by four curves – a ceiling curve, floor curve, and two additional curves inducing the two gates.

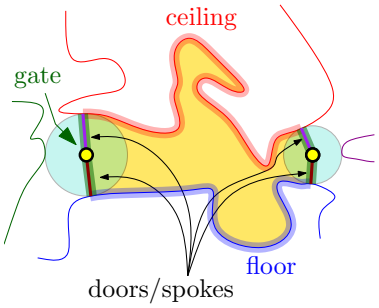

Figure 3.2: A corridor and its doors.

Here is an informal argument why corridors indeed have the above described structure. If there is any other curve, except the floor and the ceiling, that interacts with the interior of the corridor, then this curve would break the corridor into smaller sub-corridors, as there would be a medial axis vertex corresponding to three curves in the interior of the corridor, which is impossible.

3.1.2. Corridor decomposition for a subset of the curves

For a subset $\Psi \subseteq \Gamma$, of total complexity $\mathtt{t}$, one can apply a similar construction. Again, compute the medial axis of the curves of Ψ, by computing, in $O(\mathtt{t} \log \mathtt{t})$ time, the Voronoi diagram of the segments used by the curves [For87], and extracting the medial axis (it is now a planar graph instead of a tree). Now, introducing the spokes of the active vertices, results in a decomposition into corridors. For technical reasons, it is convenient to add a large bounding box, and restrict the construction to this domain, treating this ***frame*** as yet another input curve. Figure 3.4 depicts one such corridor decomposition.

Let $\mathcal{C}(\Psi)$ denote this resulting decomposition into corridors.

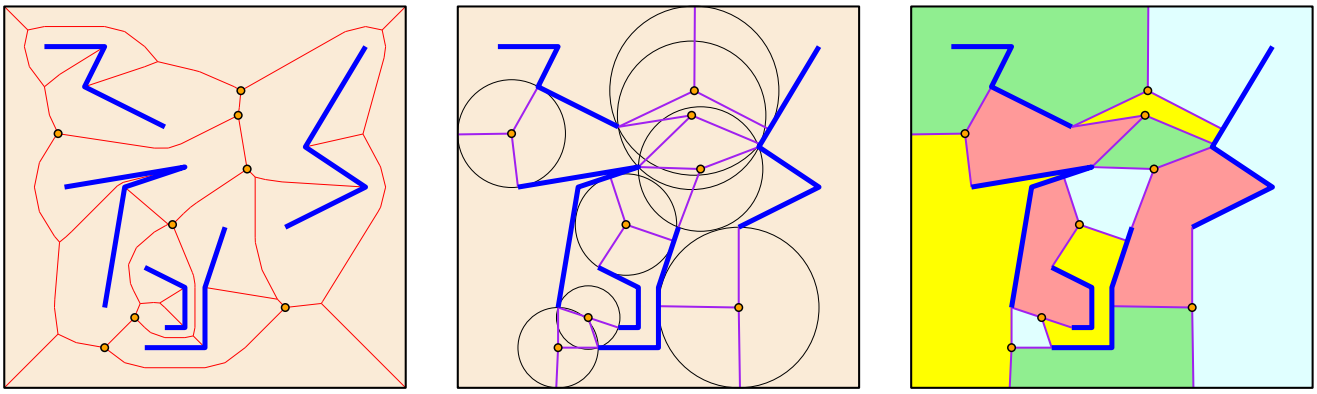

(i) (ii) (iii)

Figure 3.4: Corridor decomposition for disjoint curves: (i) Input curves, the medial axis, and active vertices, (ii) the critical circles, and their spokes, and (iii) the resulting corridor decomposition.

3.1.3. Properties of the resulting decomposition

Every corridor in the resulting decomposition $\mathcal{C}(\Psi)$ is defined by a constant number of input curves. Specifically, consider the set of all possible corridors; that is $\mathcal{F} = \bigcup_{\Upsilon \subseteq \Gamma} \mathcal{C}(\Upsilon)$. Next, consider any corridor $C \in \mathcal{F}$, then there is a unique ***defining set*** $D(C) \subseteq \Psi$ (of at most 4 curves). Similarly, such a corridor has ***stopping set*** (or ***conflict list***) of C, denoted by $K(C)$. Here, the defining set of a corridor, as depicted by Figure 3.2, is formed by the floor and the ceiling curves of the corridor, and the two external curves defining the two gates. The stopping set contains all the curves that intersect the interior of the corridor, or the interior of the two disks that define the two gates.

Consider any subset $\mathcal{S} \subseteq \Gamma$. It is easy to verify that the following two conditions hold:

(i) For any $C \in \mathcal{C}(\mathcal{S})$, we have $D(C) \subseteq \mathcal{S}$ and $\mathcal{S} \cap K(C) = \emptyset$.
(ii) If $D(C) \subseteq \mathcal{S}$ and $K(C) \cap \mathcal{S} = \emptyset$, then $C \in \mathcal{C}(\mathcal{S})$.

Namely, the corridor decomposition complies with the technique of Clarkson-Shor (see Section 2.1.2 and Definition 2.5).

3.2. Computing a specific corridor

Let p be a point in the plane, and let Γ be a set of $\overline{n}$ interior disjoint curves (stored in a read only memory), where each curve is of complexity m. Let n be the total complexity of these curves (we assume that $n = \Theta(m\,\overline{n})$). Our purpose here is to compute the corridor $C \in \mathcal{C}(\Gamma)$ that contains p. Formally, for a subset $\Psi \subseteq \Gamma$, we define the function $w(\Psi)$, to be the defining set of the corridor $C \in \mathcal{C}(\Psi)$ that contains p. Note, that such a defining set has cardinality at most $\delta = 4$.

3.2.1. Basic operations

We need to specify how to implement the two basic operations:

(A) (**Basis computation**) Given a set of $O(1)$ curves, we compute their medial axis, and extract the corridor containing p. This takes $O(m \log m)$ time, and uses $O(m)$ work space.
(B) (**Violation test**) Given a corridor C, and a curve σ, both of complexity $O(m)$, we can check if σ violates the corridor by checking if an arbitrary vertex of σ is contained in C (this takes $O(m)$ time to check), and then check in $O(m)$ time, if any segment of σ intersects the two gates of C. More precisely, each gate has an associated empty disk with it. The corridor is

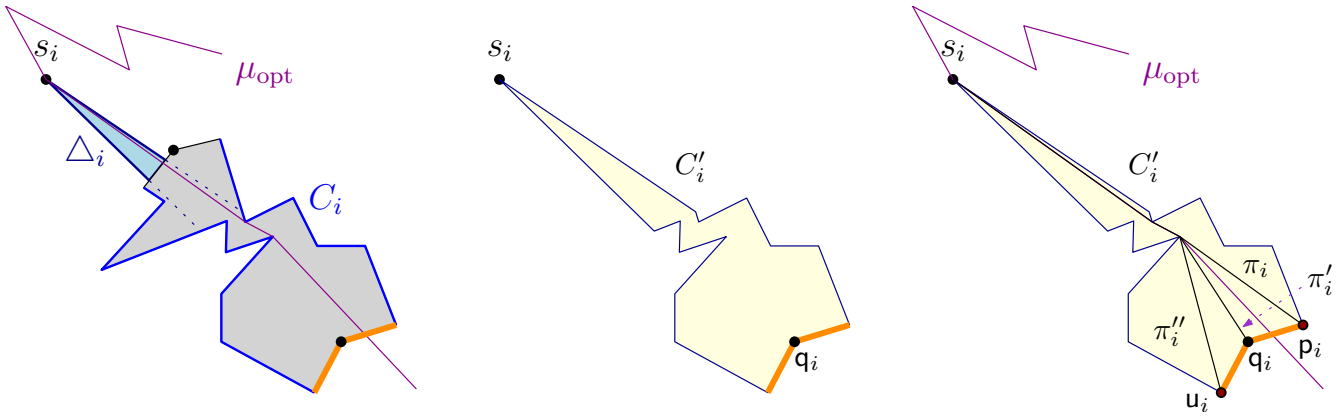

(i) (ii) (iii)

Figure 4.1: (i) The state in the beginning of the ith iteration. (ii) The clipped polygon C'_i. (iii) The funnel created by the shortest paths from s_i to the two spoke endpoints.

violated if the curve σ intersect these two empty disks associated with the two gates of the corridor, see Figure 3.2. Overall, this takes $O(m)$ time.

We thus got the following result.

Lemma 3.2. *Given a polygon P with n vertices, stored in read only memory, and let m be a parameter. Let Γ be the set of $\overline{n}$ curves resulting from breaking P into polygonal curves each with m vertices, as described in Definition 3.1. Then, given a query point p inside P, one can compute, in $O(n + m \log m \log^4(n/m))$ expected time, the corridor of $\mathcal{C}(\Gamma)$ that contains p. This algorithm uses $O(m)$ work space.*

Proof: As described above, this problem is a point location problem in a canonical decomposition; that is, a violator space problem. Plugging it into Theorem 2.18, makes in expectation $O(\overline{n})$ violation tests, each one takes $O(m)$ time, where $\overline{n} = \lfloor (n-1)/m \rfloor + 1 = O(n/m)$. The algorithm performs in expectation $O(\log^4 \overline{n})$ basis computations, each one takes $O(m \log m)$ time. Thus, expected running time is $O(\overline{n} \cdot m + m \log m \log^4 \overline{n})$ time. As for the space required, observe that a corridor can be described using $O(1)$ space.

The algorithm needs $O(m)$ work space to implement the basis computation operation, as all other portions of the algorithm require only constant additional work space. ∎

We comment that the problem of Lemma 3.2 can also be solved as an LP-type problem. We refer the interested reader to Appendix B.

4. Shortest path in a polygon in sublinear space

Let P be a simple polygon with n edges in the plane, and let s and t be two points in P, where s is the ***source***, and t is the ***target***. Our purpose here is to compute the shortest path between s and t inside P. The vertices of P are stored in (say) counterclockwise order in an array stored in a read only memory. Let m be a prespecified parameter that is (roughly) the amount of additional space available for the algorithm.

4.1. Updating the shortest path through a corridor

We remind the reader that a corridor has two double *doors*, see Figure 3.2$_{\text{p13}}$. The rest of the boundary of the corridor is made out two chains from the original polygon.

Given two rays σ and σ', that share their source vertex v (which lies inside P), consider the polygon Q that starts at v, follows the ray σ till it hits the boundary of P, then trace the boundary of P in a counterclockwise direction till the first intersection of σ' with ∂P, and then back to v. The polygon $Q = P\langle\sigma, \sigma'\rangle$ is the ***clipped*** polygon. See Figure 4.1 (ii).

A ***geodesic*** is the shortest path between two points (restricted to lie inside P). Two geodesics might have a common intersection, but they can cross at most once. Locally, inside a polygon, a geodesic is a straight segment. For our algorithm, we need the following two basic operations (both can be implemented using $O(1)$ work space):

(A) **isPntIn**(p): Given a query point p, it decides if p is inside P. This is done by scanning the edges of P one by one, and counting how many edges cross the vertical ray shooting from p downward. This operation takes linear time (in the number of vertices of P). If the count is odd then p is inside P, and it is outside otherwise[7].

(B) **isInSubPoly**(p, σ, σ'): returns true if p is in the clipped polygon $P\langle\sigma, \sigma'\rangle$. This can be implemented to work in linear time and constant space, by straightforward modification of **isPntIn**.

Using vertical and horizontal rays shot from s, one can decide, in $O(n)$ time, which quadrant around s is locally used by the shortest path from s to t. Assume that this path is in the positive quadrant. It would be useful to think about geodesics starting at s as being sorted angularly. Specifically, if τ and τ' are two geodesic starting at s, then τ is to the ***left*** of τ', if the first edge of τ is counterclockwise to the first edge of τ'. If the prefix of τ and τ' is non-empty, we apply the same test to the last common point of the two paths. Let $\tau \prec \tau'$ denote that τ is to the left of τ'.

In particular, if the endpoint of the rays σ, σ' is the source vertex s, and the geodesic between s and t lies in $P\langle\sigma, \sigma'\rangle$, then given a third ray π lying between σ and σ', the shortest path between s and t in P must lie completely either in $P\langle\sigma, \pi\rangle$ or $P\langle\pi, \sigma'\rangle$, and this can be tested by a single call to **isInSubPoly** for checking if t is in $P\langle\pi, \sigma'\rangle$.

4.1.1. Limiting the search space

Lemma 4.1. *Let P, s and t be as above, and μ be the shortest path from s to t in P. Let pq be a given edge, which is the last edge in the shortest path τ from s to q, where q is in P. Then, one can decide in $O(n)$ time, and using $O(1)$ space, if $\mu \prec \tau$, where n is the number of vertices of P.*

Proof: If $p = s$, then this can be determined by shooting two rays, one in the direction of $s \to q$ and the other in the opposite direction, where $s \to q$ denotes the vector $q - s$. Then a single call to **isInSubPoly** resolves the issue.

Otherwise, p must be a vertex of P. Consider the ray σ starting at p (or q – that does not matter) in the direction of $p \to q$. Compute, in linear time, the first intersection of this ray with the boundary of P, and let u be this point. Clearly, pu connects two points on the boundary of P, and it splits P into two polygons. One can now determine, in linear time, which of these two polygons contains t, thus resolving the issue. ∎

[7] This is the definition of inside/outside of a polygon by the Jordan curve theorem.

4.1.2. Walking through a corridor

In the beginning of the ith iteration of the algorithm it would maintains the following quantities (depicted in Figure 4.1 (i)):

(A) s_i: the current source (it lies on the shortest path μ_{opt} between s and t).
(B) C_i: The current corridor.
(C) $\triangle_i$: A triangle having s_i as one of its vertices, and its two other vertices lie on a spoke of C_i. The shortest path μ_{opt} passes through s_i, and enters C_i through the base of $\triangle_i$, and then exits the corridor through one of its "exit" spokes.

The task at hand is to trace the shortest path through C_i, in order to compute where the shortest path leaves the corridor.

Lemma 4.2. *Tracing the shortest path μ_{opt} through a single corridor takes $O\bigl(n \log m + m \log m \log^4 \overline{n}\bigr)$ expected time, using $O(m)$ space.*

Proof: We use the above notation. The algorithm glues together $\triangle_i$ to C_i to get a new polygon. Next, it clips the new polygon by extending the two edges of $\triangle_i$ from s_i. Let C'_i denote the resulting polygon, depicted in Figure 4.1 (ii). Let the three vertices of C'_i forming the two "exit" spokes be p_i, q_i, u_i. Next, the algorithm computes the shortest path from s_i to the three vertices p_i, q_i, u_i inside C'_i, and let π_i, π'_i, π''_i be these paths, respectively (this takes $O(|C'_i|) = O(m)$ time [GH89] after linear time triangulation of the polygon [Cha91; AGR01]). Using Lemma 4.1 the algorithm decides if $\pi_i \prec \mu_{\text{opt}} \prec \pi'_i$ or $\pi'_i \prec \mu_{\text{opt}} \prec \pi''_i$. We refer to a prefix path (that is part of the desired shortest path) followed by the two concave chains as a ***funnel*** – see Figure 4.1 (iii) and Figure 4.2 for an example.

Assume that $\pi_i \prec \mu_{\text{opt}} \prec \pi'_i$, and let F_i be the funnel created by these two shortest paths, where $p_i q_i$ is the base of the funnel. If the space bounded by the funnel is a triangle, then the algorithm sets its top vertex as s_{i+1}, the funnel triangle is $\triangle_{i+1}$, and the algorithm computes the corridor on the other side of $p_i q_i$ using the algorithm of Lemma 3.2 (by picking any point on $p_i q_i$ as the query for the point location operation), set it as C_{i+1}, and continues the execution of the algorithm to the next iteration.

The challenge is handling funnel chains that are "complicated" concave polygons (with at most $O(m)$ vertices), see Figure 4.2. As long as the funnel F_i is not a triangle, pick a middle edge e on one side of the funnel, and extend it till it hits the edge $p_i p_{i+1}$, at a point x. This breaks F_i into two funnels, and using the algorithm of Lemma 4.1 on the edge e, decide which of these two funnels contains the shortest path μ_{opt}, and replace F_i by this funnel. Repeat this process till F_i becomes a triangle. Once this happens, the algorithm continues to the next iteration as described above. Clearly, this funnel "reduction" requires $O(\log m)$ calls to the algorithm of Lemma 4.1.

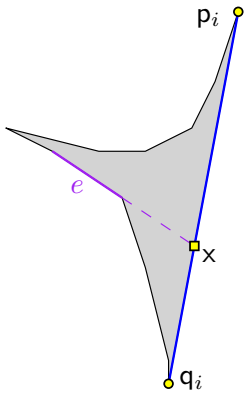

Figure 4.2: Funnel reduction.

Note, that the algorithm "forgets" the portion of the funnel that is common to both paths as it moves from C_i to C_{i+1}. This polygonal path is a part of the shortest path μ_{opt} computed by the algorithm, and it can be output at this stage, before moving to the next corridor C_{i+1}.

In the end of the iteration, this algorithm computes the next corridor C_{i+1} by calling the algorithm of Lemma 3.2. ∎

4.2. The algorithm

The overall algorithm works by first computing the corridor C_1 containing the source $s_1 = s$ using Lemma 3.2. The algorithm now iteratively applies Lemma 4.2 till arriving to the corridor containing t,

where the remaining shortest path can be readily computed. Since every corridor gets visited only once by this walk, we get the following result.

Theorem 4.3. *Given a simple polygon P with n vertices (stored in a read only memory), a start vertex s, a target vertex t, and a space parameter m, one can compute the length of the shortest path from s to t (and output it), using $O(m)$ additional space, in $O(n^2/m)$ expected time, if $m = O(n/\log^2 n)$. Otherwise, the expected running time is $O\big(n^2/m + n \log m \log^4(n/m)\big)$.*

Proof: The algorithm is described above, and let $\overline{n} = \lfloor (n-1)/m \rfloor + 1$. There are $O(\overline{n})$ corridors, and this bounds the number of iterations of the algorithm. As such, the overall expected running time is $O\big(\overline{n}\big(n \log m + m \log m \log^4 \overline{n}\big)\big) = O\big((n^2/m) \log m + n \log m \log^4 \overline{n}\big)$.

To get a better running time, observe that the extra log factor (on the first term), is rising out of the funnel reduction $O(\log m)$ queries inside each corridor, done in the algorithm of Lemma 4.2. If instead of reducing a funnel all the way to constant size, we reduce it to have say, at most $\lceil m/4 \rceil$ edges (triggered by the event that the funnel has at least $m/2$ edges), then at each invocation of Lemma 4.2, only a constant number of such queries would be performed. One has to adapt the algorithm such that instead of a triangle entering a new corridor, it is a funnel. The adaptation is straightforward, and we omit the easy details. The improved running time is $O\big(n^2/m + n \log m \log^4 \overline{n}\big)$. ∎

5. Epilogue

In this paper, we showed the following:

(A) Violator spaces are the natural way to phrase the problem of one-shot point location query in a canonical decomposition.
(B) A variant of Seidel's algorithm solves Violator spaces in expected linear time.
(C) This algorithm can be modified to use only $O(1)$ work space, by using pseudo-random sequences.
(D) Backward analysis works if one uses pseudo-random sequences instead of true random permutations (Lemma 2.16).
(E) Last, but not least, shortest path in a polygon can be computed in $O(n^2/m)$ expected time, using $O(m)$ work space. This is done by breaking a polygon into $O(n/m)$ subpolygons of size $O(m)$, and walking through these corridors, spending $O(n)$ time in each corridor.

While deriving some of the above results required quite a bit of technical mess[8] (specifically (C) and (D) above), conceptually the basic ideas themselves are simple and natural.

The most interesting open problem remaining from our work, is whether one can improve the running time for computing the shortest path in a polygon with $O(m)$ space to be faster than $O(n^2/m)$.

Acknowledgments

The author became aware of the low-space shortest path problem during Tetsuo Asano talk in the Workshop in honor of his 65th birthday during SoCG 2014. The author thanks him for the talk, and the subsequent discussions. The author also thanks Pankaj Agarwal, Chandra Chekuri, Jeff Erickson and Bernd Gärtner for useful discussions. The author also thanks the anonymous referees, for both the conference and journal versions, for their detailed comments.

[8] The author would like to use this historical opportunity to apologize both to the reader and to himself for this mess.

References

[AGR01] N. M. Amato, M. T. Goodrich, and E. A. Ramos. *A randomized algorithm for triangulating a simple polygon in linear time*. *Discrete Comput. Geom.*26.(2): 245–265, 2001.

[Asa+11] T. Asano, W. Mulzer, G. Rote, and Y. Wang. *Constant-work-space algorithms for geometric problems*. *J. Comput. Geom.*2.(1): 46–68, 2011.

[Asa+13] T. Asano, K. Buchin, M. Buchin, M. Korman, W. Mulzer, G. Rote, and A. Schulz. *Memory-constrained algorithms for simple polygons*. *Comput. Geom. Theory Appl.*46.(8): 959–969, 2013.

[Asa+14] T. Asano, K. Buchin, M. Buchin, M. Korman, W. Mulzer, G. Rote, and A. Schulz. *Reprint of: Memory-constrained algorithms for simple polygons*. *Comput. Geom. Theory Appl.*47.(3): 469–479, 2014.

[Asa12] T. Asano. *In-place algorithm for erasing a connected component in a binary image*. *Theo. Comp. Sci.*50.(1): 111–123, 2012.

[Ber+08] M. Berg, O. Cheong, M. Kreveld, and M. H. Overmars. *Computational Geometry: Algorithms and Applications*. 3rd. Springer-Verlag, 2008.

[BG11] Y. Brise and B. Gärtner. *Clarkson's algorithm for violator spaces*. *Comput. Geom. Theory Appl.*44.(2): 70–81, 2011.

[Blu67] H. Blum. *A transformation for extracting new descriptors of shape*. *Models for the Perception of Speech and Visual Form*. Ed. by W. Wathen-Dunn. MIT Press, 1967, pp. 362–380.

[CC07] T. M. Chan and E. Y. Chen. *Multi-pass geometric algorithms*. *Discrete Comput. Geom.*37.(1): 79–102, 2007.

[Cha91] B. Chazelle. *Triangulating a simple polygon in linear time*. *Discrete Comput. Geom.*6: 485–524, 1991.

[CHR15] H.-C. Chang, S. Har-Peled, and B. Raichel. *From proximity to utility: A Voronoi partition of Pareto optima*. *Proc. 31st Int. Annu. Sympos. Comput. Geom.* (SoCG), vol. 34. 2015, 689–703, 2015.

[Cla88] K. L. Clarkson. *Applications of random sampling in computational geometry, II*. *Proc. 4th Annu. Sympos. Comput. Geom.* (SoCG), 1988, 1–11, 1988.

[Cla95] K. L. Clarkson. *Las Vegas algorithms for linear and integer programming*. *J. Assoc. Comput. Mach.*42: 488–499, 1995.

[CLM05] B. Chazelle, D. Liu, and A. Magen. *Sublinear geometric algorithms*. *SIAM J. Comput.*35.(3): 627–646, 2005.

[CM96] B. Chazelle and J. Matoušek. *On linear-time deterministic algorithms for optimization problems in fixed dimension*. *J. Algorithms.*21: 579–597, 1996.

[CS89] K. L. Clarkson and P. W. Shor. *Applications of random sampling in computational geometry, II*. *Discrete Comput. Geom.*4: 387–421, 1989.

[CSW99] F. Y. L. Chin, J. Snoeyink, and C. A. Wang. *Finding the medial axis of a simple polygon in linear time*. *Discrete Comput. Geom.*21.(3): 405–420, 1999.

[For87] S. J. Fortune. *A sweepline algorithm for Voronoi diagrams*. *Algorithmica.*2: 153–174, 1987.

[GH89] L. J. Guibas and J. Hershberger. *Optimal shortest path queries in a simple polygon*. *J. Comput. Sys. Sci.*39.(2): 126–152, 1989.

[Gär+08] B. Gärtner, J. Matoušek, L. Rüst, and P. Škovroň. *Violator spaces: Structure and algorithms*. *Discrete Appl. Math.*156.(11): 2124–2141, 2008.

[Gär+06] B. Gärtner, J. Matoušek, L. Rüst, and P. Škovroň. *Violator spaces: Structure and algorithms*. *Proc. 14th Annu. Euro. Sympos. Alg.* (ESA), 2006, 387–398, 2006.

[Har15] S. Har-Peled. *Shortest path in a polygon using sublinear space*. *Proc. 31st Int. Annu. Sympos. Comput. Geom.* (SoCG), vol. 34. 2015, 111–125, 2015.

[Ind01] P. Indyk. *A small approximately min-wise independent family of hash functions*. *J. Algorithms.*38.(1): 84–90, 2001.

[Ind06] P. Indyk. *Stable distributions, pseudorandom generators, embeddings, and data stream computation*. *J. Assoc. Comput. Mach.*53.(3): 307–323, 2006.

[KM88] S. Kapoor and S. N. Maheshwari. *Efficient algorithms for Euclidean shortest path and visibility problems with polygonal obstacles*. *Proc. 4th Annu. Sympos. Comput. Geom.* (SoCG), 1988, 172–182, 1988.

[LP84] D. T. Lee and F. P. Preparata. *Euclidean shortest paths in the presence of rectilinear barriers*. *Networks.*14: 393–410, 1984.

[Meg84] N. Megiddo. *Linear programming in linear time when the dimension is fixed*. *J. Assoc. Comput. Mach.*31: 114–127, 1984.

[MN98] J. Matoušek and J. Nešetřil. *Invitation to Discrete Mathematics*. Oxford Univ Press, 1998.

[MR96] I. J. Munro and V. Raman. *Selection from read-only memory and sorting with minimum data movement*. *Theo. Comp. Sci.*165.(2): 311–323, 1996.

[Mul94] K. Mulmuley. *Computational Geometry: An Introduction Through Randomized Algorithms*. Prentice Hall, 1994.

[Rüs07] L. Y. Rüst. *The P-Matrix Linear Complementarity Problem – Generalizations and Specializations*. Diss. ETH No. 17387. PhD thesis. ETH, 2007.

[Sei91] R. Seidel. *Small-dimensional linear programming and convex hulls made easy*. *Discrete Comput. Geom.*6: 423–434, 1991.

[SW92] M. Sharir and E. Welzl. *A combinatorial bound for linear programming and related problems*. *Proc. 9th Annul. Sympos. Theoret. Asp. Comput. Sci.* (STACS), vol. 577. 1992, 569–579, 1992.

[Har11] S. Har-Peled. *Geometric Approximation Algorithms*. Vol. 173. Mathematical Surveys and Monographs. Amer. Math. Soc., 2011.

[Har14] S. Har-Peled. *Quasi-polynomial time approximation scheme for sparse subsets of polygons*. *Proc. 30th Annu. Sympos. Comput. Geom.* (SoCG), 2014, 120–129, 2014.

[Ško07] P. Škovroň. *Abstract models of optimization problems*. `http://kam.mff.cuni.cz/~xofon/thesis/diplomka.pdf`. PhD thesis. Charles University, 2007.

A. Backward analysis for pseudo-random sequences

We prove the probabilistic bounds we need from scratch. We emphasize that this is done so that the presentation is self contained. Our estimates and arguments are inspired by Mulmuley's [Mul94, Chapter 10], although we are dealing with somewhat different events. In particular, there is nothing new in Appendix A.1, and relatively little new in Appendix A.2. Finally, Appendix A.3 proves the new bounds we need.

Interestingly, Indyk [Ind01] used arguments in a similar spirit to bound a different event, related to the probability of the ith element in a pseudo-random "permutation" to be a minimum of all the elements seen so far. This does not seem to have a direct connection to the analysis here.

A.1. Some probability stuff

We need the following lemma [Mul94, Theorem A.2.1] and provide a somewhat more precise bounds on the constants involved, than the reference mentioned.

Restatement of Lemma 2.13. *Let $X_1, \ldots, X_n$ be n random indicator variables that are 2ψ-wise independent, where $p = \mathbf{Pr}[X_i = 1] \geq 1/n$, for all i. Let $Y = \sum_i X_i$, and let $\mu = \mathbf{E}[Y] = pn$. Then, we have that $\mathbf{Pr}\big[Y = 0\big] \leq \mathbf{Pr}\big[|Y - \mu| \geq \mu\big] \leq \big(4\psi^2/\mu\big)^{\psi}/\sqrt{2\psi}$.*

Proof: For $i > 1$, and any j, let

$$\alpha_i = \mathbf{E}\Big[(X_j - p)^i\Big] = p(1-p)^i + (1-p)(-p)^i \leq p(1-p)\Big((1-p)^{i-1} + p^{i-1}\Big) \leq p(1-p) \leq p.$$

Similarly, we have that $\alpha_1 = \mathbf{E}[X_j - p] = \mathbf{E}[X_j] - p = p - p = 0$.

Consider the variable $Z = \Big(\sum_{j=1}^{n}(X_j - p)\Big)^{2\psi}$, and its expectation. Let $\mathcal{I}$ be the set of all 2ψ tuples $(i_1, \ldots, i_{2\psi}) \in [\![n]\!]^{2\psi}$, such that $i_1 \leq i_2 \leq \ldots \leq i_{2\psi}$. We have by linearity of expectations that $\mathbf{E}\big[Z\big] = \sum_{(i_1,\ldots,i_{2\psi})\in\mathcal{I}} \mathbf{E}\Big[\prod_{j=1}^{2\psi}\big(X_{i_j} - p\big)\Big]$. Now, any tuple with $(i_1, \ldots, i_{2\psi})$ such that an index, say i_ℓ, appears exactly once, contributes 0 to this summation, since then

$$\mathbf{E}\Big[\prod\nolimits_{j=1}^{2\psi}\big(X_{i_j} - p\big)\Big] = \mathbf{E}\Big[X_{i_\ell} - p\Big]\,\mathbf{E}\Big[\prod\nolimits_{j\in[\![2\psi]\!]\setminus\{\ell\}}\big(X_{i_j} - p\big)\Big] = 0,$$

as $\alpha_1 = 0$, and as the variables are 2ψ-wise independent. So, consider a tuple $I \in \mathcal{I}$, where every index appears at least twice, and there are $\nu \leq \psi$ distinct indices. In particular, if such a term involves variables $i_1, \ldots, i_\nu$, with multiplicities $n_1, \ldots, n_\nu$ (all at least 2), then

$$\mathbf{E}\Big[\prod\nolimits_{j=1}^{\nu}\big(X_{i_j} - p\big)^{n_j}\Big] = \prod\nolimits_{j=1}^{\nu}\mathbf{E}\Big[\big(X_{i_j} - p\big)^{n_j}\Big] \leq \prod\nolimits_{j=1}^{\nu}\alpha_{n_j} \leq p^{\nu},$$

using the 2ψ-wise independence.

We want to bound the total sum of the coefficients of all such monomials in the polynomial Z that have at most ν distinct variables. To this end, to specify how we generated such a monomial, we need to specify 2ψ integers $i_1, \ldots, i_{2\psi}$. We first scan such a sequence and write down the ν unique values encountered, in the order they are being encountered. There are $\leq n^\nu$ such "seen first" sequences. Now, for every index i_t, we need to specify if it is new or not. There are $\binom{2\psi}{\nu}$ such choices. Finally, for each number in this sequence which is a repetition, we need to specify which of the previously (at most

ν numbers) seen before it is. There are $\nu^{2\psi-\nu}$ such possibilities. As such, the total sum of all such coefficients in Z is bounded by $\binom{2\psi}{\nu} n^\nu \nu^{2\psi-\nu}$. Every such monomial contributes p^ν to $\mathbf{E}[Z]$. As such, we have

$$\mathbf{E}\Big[Z\Big] \leq 0 + \sum_{\nu=1}^{\psi} \binom{2\psi}{\nu} n^\nu \nu^{2\psi-\nu} p^\nu \leq (pn)^\psi \sum_{\nu=1}^{\psi} \binom{2\psi}{\nu} \psi^{2\psi-\nu} \leq \frac{2^{2\psi}}{\sqrt{2\psi}} \psi^{2\psi} (pn)^\psi = \frac{(4\psi^2 pn)^\psi}{\sqrt{2\psi}},$$

since $pn \geq 1$, and $\binom{2\psi}{\nu} \leq \binom{2\psi}{\psi} \leq 2^\psi/\sqrt{2\psi}$ (see [MN98, Section 2.6]).

By Markov's inequality, we have that

$$\mathbf{Pr}\Big[|Y-\mu| \geq \mu\Big] = \mathbf{Pr}\Big[(Y-\mu)^{2\psi} \geq \mu^{2\psi}\Big] \leq \frac{\mathbf{E}\Big[(Y-\mu)^{2\psi}\Big]}{\mu^{2\psi}} = \frac{\mathbf{E}\big[Z\big]}{\mu^{2\psi}} \leq \frac{(4\psi^2)^\psi \mu^\psi}{\sqrt{2\psi}\mu^{2\psi}},$$

implying the claim. ∎

A.2. The probability of a specific basis to have a conflict at iteration i

The following lemma bounds the probability that a specific basis B is defined in the prefix of the $i-1$ sampled elements, and the first element violating it is in position i. This probability is of course affected by the size of its conflict list L.

Lemma A.1. *Let $X_1, \ldots X_i$ be a sequence of variables that are uniformly distributed over $[\![n]\!]$, that are ϕ-wise independent, where $\phi > \delta$ is a sufficiently large constant. Let $B, L \subseteq [\![n]\!]$ be two disjoint sets of size δ and $k = t(n/i)$, respectively, where δ is a constant, $i > 2\delta$, and $t > 0$ is arbitrary. Let $\mathcal{E}$ be the following event:*

(I) $B \subseteq \{X_1, \ldots, X_{i-1}\}$. (II) $L \cap \{X_1, \ldots, X_{i-1}\} = \emptyset$. (III) $X_i \in L$.

Then, for $t > 1$, we have $\mathbf{Pr}\big[\mathcal{E}\big] = O\Big(\frac{1}{i}\big(\frac{i}{n}\big)^\delta \frac{1}{t^\psi}\Big)$, where $\psi = (\phi - \delta - 1)/2 - 1$.

For $t \leq 1$, we have $O\Big(\frac{1}{i}\big(\frac{i}{n}\big)^\delta\Big)$, as long as $\phi \geq \delta$.

Proof: Fix the numbering of the elements of B as $B = \{b_1, \ldots, b_\delta\}$, and let $\mathcal{I}$ be the set of tuples of δ distinct indices in $[\![i-1]\!]^\delta$. For such a tuple $I = (i_1, \ldots, i_\delta)$, let $\mathcal{E}(I)$ be the event that $X_{i_1} = b_1, \ldots, X_{i_\delta} = b_\delta$, and $X_i \in L$. We have that

$$\mathbf{Pr}\Big[\mathcal{E}(I)\Big] = \mathbf{Pr}\Big[X_i \in L\Big] \prod_{j=1}^{\delta} \mathbf{Pr}\Big[X_{i_j} = b_j\Big] = \frac{1}{n^\delta} \cdot \frac{k}{n} = \frac{k}{n^{\delta+1}},$$

as the variables $X_1, \ldots, X_n$ are ϕ-wise independent, and $\phi > \delta + 1$. Let $\mathcal{F}(I)$ be that event that none of the variables of $\mathcal{X} = \{X_1, \ldots, X_{i-1}\}$ are in L. In particular, consider the $m = i - 1 - \delta$ variables in $\mathcal{X}$ that have an index in $[\![i-1]\!] \setminus I$, and let X'_j denote the jth such variable. Let Y_j be an indicator variable that is one if $X'_j \in L$. The variables $X'_1, \ldots, X'_m$ are (at worst) $(\phi - \delta - 1)$-wise independent (since, conceptually, we fixed the value of the variables of I), and so are the indicator variables $Y_1, \ldots, Y_m$. For any j, we have that $p = \mathbf{E}[Y_j] = \mathbf{Pr}\big[X'_j \in L\big] = k/n$. As such, for $\psi = (\phi - \delta - 1)/2$, we have that

$$\begin{aligned}
\mathbf{Pr}\Big[\mathcal{F}(I) \cap \mathcal{E}(I)\Big] &= \mathbf{Pr}\Big[\mathcal{F}(I) \,\Big|\, \mathcal{E}(I)\Big] \mathbf{Pr}\Big[\mathcal{E}(I)\Big] = \mathbf{Pr}\Big[\textstyle\sum_j Y_j = 0\Big] \frac{k}{n^{\delta+1}} \\
&= O\left(\left(\frac{1}{pm}\right)^{(\phi-\delta-1)/2} \frac{k}{n^{\delta+1}}\right) = O\left(\left(\frac{n}{k(i-1-\delta)}\right)^\psi \frac{k}{n^{\delta+1}}\right) \\
&= O\left(\left(\frac{n}{ki}\right)^\psi \frac{(tn/i)}{n^{\delta+1}}\right) = O\left(\frac{1}{t^\psi} \cdot \frac{t}{in^\delta}\right) = O\left(\frac{1}{it^{\psi-1}n^\delta}\right),
\end{aligned}$$

by Lemma 2.13, as $i > 2\delta$, ϕ is constant, and $k = tn/i$. There are at most i^δ choices for the tuple I, and there are $\delta!$ different orderings of the elements of B. As such, the desired probability is bounded by $O\left(i^\delta \delta! \frac{1}{it^{\psi-1}n^\delta}\right) = O\left(\frac{1}{i}\left(\frac{i}{n}\right)^\delta \frac{1}{t^{\psi-1}}\right)$, as δ is a constant.

The bound for $t \leq 1$, follows by using $\mathcal{E}(I)$ instead of $\mathcal{F}(I) \cap \mathcal{E}(I)$ in the above analysis. ∎

A.3. Back to backward analysis

Consider a fixed violator space $\mathcal{V} = (H, \mathrm{cl})$ with $n = |H|$ constraints, and combinatorial dimension δ. Let $\mathcal{B}$ be the set of bases of $\mathcal{V}$. Let $\mathcal{X} = X_1, \ldots, X_n$ be the sequence of constraints generated by the pseudo-random generator, that is ϕ-wise independent (say, somewhat arbitrarily, $\phi = 6\delta + 10$). In the following, the *ith prefix* of $\mathcal{X}$ is $\mathcal{X}_i = \langle X_1, \ldots, X_i \rangle$.

Lemma A.2. *Let $\mathcal{B}_{\leq k}$ be the set of all the bases of $\mathcal{V}$ that have a conflict list of size at most k. Then, we have that $|\mathcal{B}_{\leq k}| = O(k^\delta)$.*

Proof: Follows readily from the Clarkson-Shor [Cla88; CS89] technique. Pick every constraint into R with probability $1/k$ into a sample R. The probability of a basis with conflict list of size at most k to survive this purge and be the basis of R is $\Omega(1/k^\delta)$. Since basis(R) is a single basis, the result readily follows. ∎

Lemma A.3. *Assume that $\phi \geq \delta$. For $i > 2\delta$, let $\rho_0(i)$ be the probability that X_i violates $B = \mathrm{basis}(X_1, \ldots, X_{i-1}) \in \mathcal{B}_{\leq k}$, where $k = n/i$. Then $\rho_0(i) = O(1/i)$.*

Proof: For a fixed basis $B \in \mathcal{B}_{\leq k}$ the desired probability, for B being the basis, is $\rho_0(i, B) = O\left(\frac{1}{i}\left(\frac{i}{n}\right)^\delta\right)$, by Lemma A.1 (for $t \leq 1$). By Lemma A.2, we have $|\mathcal{B}_{\leq k}| = O((n/i)^\delta)$. As such, the desired probability is bounded by $O\left((1/i)(i/n)^\delta (n/i)^\delta\right) = O(1/i)$. ∎

A basis B is ***t-heavy*** for $\mathcal{X}_i$, if $\frac{|\mathrm{cl}(B)|}{n/i} \in (t-1, t]$.

Lemma A.4. *Assume that ϕ is a constant and $\phi \geq 3(\delta+3)$. For $i > 2\delta$, let $\rho_t(i)$ be the probability that X_i violates $B = \mathrm{basis}(X_1, \ldots, X_{i-1}) \in \mathcal{B}$, and B is t-heavy, for $t > 2$. Then $\rho_t(i) = O(1/(it^2))$.*

Proof: Arguing as in Lemma A.3, there are $O\left((tn/i)^\delta\right)$ t-heavy bases, and by Lemma A.1, the probability of the desired event for a specific basis B', with $|\mathrm{cl}(B')| \geq (t-1)n/i$, is bounded by $O\left(\frac{1}{i}\left(\frac{i}{n}\right)^\delta \frac{1}{t^\psi}\right)$. As such, the desired probability is bounded by $O\left(\left(\frac{tn}{i}\right)^\delta \frac{1}{i}\left(\frac{i}{n}\right)^\delta \frac{1}{t^\psi}\right) = O\left(\frac{1}{it^2}\right)$, if $\psi \geq \delta + 2$, which holds for $\phi \geq 3(\delta+3)$. ∎

Restatement of Lemma 2.16. *Consider a violator space $\mathcal{V} = (H, \mathrm{cl})$ with $n = |H|$, and combinatorial dimension δ. Let $\mathcal{X} = X_1, X_2, \ldots,$ be a random sequence of constraints of H generated by a ϕ-wise independent distribution (with each X_i having a uniform distribution), where $\phi > 6\delta + 9$ is a constant. Then, for $i > 2\delta$, the probability that X_i violates $B = \mathrm{basis}(X_1, \ldots, X_{i-1})$ is $O(1/i)$.*

Proof: Plugging in the bounds of Lemma A.3 and Lemma A.4, we have that the desired probability is bounded by $\rho_0(i) + \sum_{t=2}^{\infty} \rho_t(i) = O\left(1/i + \sum_{t=2}^{\infty} \frac{1}{it^2}\right) = O(1/i)$, as claimed. ∎

B. point location in corridors using LP-type solver

For the amusement of the interested reader, we briefly describe how point location in corridor decomposition can be solved as an LP-type problem. This requires only defining a "crazy" ordering on the corridors that contain our query point, and the rest follows readily. Of course, it is more elegant and natural to do this using violator spaces.

We need to define an explicit ordering on the corridors so we can use the LP-type algorithm. This is somewhat tedious. The corridor C corresponds to an edge e of the medial axis connecting the two active vertices u and x of the medial axis. Let D be the set of four inputs curves that define the corridor C (i.e., D is the *defining set* of C). The closest curve to the query point p is the ***floor*** of C. The other curve of D that lies on the other side of the medial axis edge e, is the ***ceiling*** of the corridor. Consider the closest point p' on the medial axis to p, and consider the two paths from p' to u and x. These two paths diverge at a point y. The curve in D corresponding to the active vertex for which the path diverges to the right (resp. left) at y, is the ***right*** (resp. ***left***) curve of C.

We can now define an order on two corridors $w(F), w(F')$ that contain p. Formally, we define

- If the floor of $w(F)$ is closer to p than $w(F')$, then we define $w(F) \succ w(F')$.
- Otherwise, by our general position assumption, it must be that the floors of $w(F)$ and $w(F')$ are the same. If the ceilings of $w(F)$ and $w(F')$ are different, then consider the ceiling c of $w(F \cup F')$. If c is also the ceiling of $w(F)$, then $w(F) \succ w(F')$, otherwise, c is the ceiling of $w(F')$, and $w(F') \prec w(F)$.
- Otherwise, $w(F)$ and $w(F')$ have the same floor and the same ceiling. We apply the same kind of argument as above to the left side of the corridor to resolve the order, and if this does not work, we use the right side of the corridor to resolve the ordering.

Now, it is easy to verify that computing $w(\Gamma)$ is an LP-type problem, with the combinatorial dimension being $\delta = 4$. Thus, the maximum corridor in this decomposition is the desired answer to the point location query, and this now can be solved using this ordering via an LP-type solver.